ROBOTS OR WORKERS?

A MACRO ANALYSIS OF AUTOMATION AND LABOR MARKETS

SYLVAIN LEDUC AND ZHENG LIU

Abstract. We study the implications of automation for labor market fluctuations in a DMP framework, generalized to incorporate automation decisions. If a job opening is not filled with a worker, a firm can choose to automate that position and use a robot instead of a worker to produce output. The threat of automation strengthens the firm's bargaining power against job seekers in wage negotiations, depressing equilibrium real wages in a business cycle boom. The option of automation also increases the value of a vacancy, raising the incentive for job creation, and thereby amplifying fluctuations in vacancies and unemployment relative to the standard DMP framework. Since automation improves labor productivity while muting wage increases, it implies a counter-cyclical labor income share, as observed in the data.

I. Introduction

Recent development in robotics and artificial intelligence have renewed concerns that robots could render workers redundant, possibly leading to what Keynes (1930) called "technological unemployment." The decline in the labor income share since the turn of the century and stagnant wage growth during the recovery from the Great Recession have also heightened fears that automation, by substituting robots for workers, is weakening workers' bargaining power. However, the notion of "machines replacing workers" is an over-simplification of the macroeconomic impact of automation (Autor, 2015). While robots have been increasingly adopted to perform standardized tasks previously performed by workers, new tasks for which workers have a comparative advantage are also being created.(Acemoglu and Restrepo, 2018). As a result of these different channels, the impact of automation on employment, wages, and the labor share is a priori ambiguous.

Date: June 18, 2019.

Key words and phrases. Robots, threat of automation, unemployment, wage bargaining, productivity, business cycles.

JEL classification: E32, J63, J64.

Leduc: Federal Reserve Bank of San Francisco. Email: Sylvain.Leduc@sf.frb.org. Liu: Federal Reserve Bank of San Francisco. Email: Zheng.Liu@sf.frb.org. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or of the Federal Reserve System.

In this paper, we examine the interactions between automation and the labor market by developing a general equilibrium framework with a frictional labor market and the option to automate jobs. We focus on the implications of automation for labor market fluctuations over the business cycles. In particular, we examine the mechanism through which automation drives changes in unemployment, real wages, and labor productivity. We assess the quantitative importance of the automation mechanism by estimating the model to fit quarterly U.S. time series and compare the labor market implications of the estimated model with those of a counterfactual version with no changes in automation decisions.

Our model builds on the standard Diamond-Mortensen-Pissarides (DMP) labor-search model and generalizes it to incorporate automation decisions. The model has two key features. First, firms create vacancies at a fixed cost. A vacancy is created if the realized i.i.d. draw of the vacancy creation cost is below the value of the vacancy. Thus, unlike the standard DMP model, entry is costly in our model and a vacancy carries a positive value (Leduc and Liu, 2019). This positive value is a necessary condition for automation to impact workers' effective bargaining power. In addition. costly vacancy creation induces more persistent movements in vacancies and thus better align the model with the empirical evidence.

Second, we incorporate endogenous automation decisions. To keep the model tractable, we assume that automation applies to a job position. We interpret a job position broadly as consisting of a bundle of tasks, which are ex ante identical, but a fraction of which will be automated depending on the realization of the idiosyncratic costs for automation. This approach simplifies our analysis significantly.[1] After posting a vacancy, a firm can fill the position with a worker and obtain the employment value. If the position is not filled, the firm draws an i.i.d. cost of automation. If the realized automation cost is below the expected benefit of automation, then the position is automated, in which case the firm uses a robot instead of a worker to produce output. If the unfilled position is not automated because of high cost draws, the firm keeps the position open and receives the continuation value of the vacancy, including the option to automate the position in future periods.

In the aggregate economy, newly hired workers add to the employment pool, and a fraction of the employment stock is separated in each period. Similarly, newly adopted robots add to the automation stock, which becomes obsolete over time at a constant rate. For simplicity, we assume that goods produced by workers and by robots are perfect substitutes. Thus, aggregate output is just the sum of the outputs produced by the two types of technologies.

[1]Acemoglu and Restrepo (2018) use an alternative framework to study automation. Building on the earlier work of Zeira (1998), they consider a job consisting of a continuum of tasks, a fraction of which are technologically automatable, while the other tasks need to be performed by human workers (see also Autor and Salomons (2018)).

Our model yields a few theoretical insights.

First, automation has a direct job-displacing effects since goods produced by robots are perfect substitutes for goods produced by workers. But automation has also a job-creation effect: the option of automating an unfilled job vacancy boosts the present value of a vacancy, raising the incentive for job creation. The net effect of automation on employment can be ambiguous, depending on the relative strength of the two opposing effects. Under our estimated parameters, automation amplifies unemployment fluctuations.

Second, the threat of automation dampens wage increases in a business cycle boom. Since the net value of automation is procyclical, the probability of automation increases in good economic times, raising the firm's reservation value (i.e., the value of a vacancy) in wage bargaining, and therefore muting wage increases.

Third, increased automation in a boom raises aggregate productivity, further fueling the expansion. By dampening wage increases while boosting labor productivity, automation contributes to a decline in the labor income share in a business cycle expansion.

We further show that, in a counterfactual model without automation, exogenously reducing workers' bargaining weight can also amplify unemployment fluctuations and dampen wage adjustments, similar to the automation channel in our framework. However, the responses of the labor share differ qualitatively. For instance, in response to a positive demand shock, our model predicts a decline in the real wage, an increase in labor productivity, and thus a decline in the labor share. The counterfacutal model without automation and with a low workers' bargainbing weight predicts a small increase in the real wage, no change in labor productivity, and thus an increase in the labor share. This difference reflects the threat of automation on wage bargaining and the endogenous productivity enhancement from automation.

To assess the quantitative importance of our mechanism, we estimate the model to fit quarterly U.S. time series data. These time series include unemployment, vacancies, real wage growth, and non-farm business sector labor productivity growth, with a sample ranging from 1964:Q2 to 2018:Q4. To fit these four time series, we assume four shocks in our model, including a discount factor shock, a neutral technology shock, an automation-specific shock, and a job separation shock. We find that matching the movements in labor productivity is an important disciplining device on our endogenous automation mechanism, given the important decline in productivity growth since the mid-2000s. In addition, note that since we fit the growth in real wages and productivity, we necessarily fit the growth in the labor share by construction over the estimation period.

We find that the option to automate plays an important role behind the decline in the labor share, particularly during the recovery from the Great Recession. In a counterfactual exercise that keeps the probability of automation constant to its steady state value, and that thus abstracts from endogenous responses in automation, our model predicts a roughly constant labor share in contrast to its sharp decline following the Great Recession. Absent automation and its threat, wages would have been higher during this period as well, while productivity would have been even more muted. Automation's negative impact on the labor share is in line with the evidence in Autor and Salomons (2018). Using data on 28 industries across 18 OECD countries between 1970 and 2007, they find that while automation increases total hours worked it nonetheless lowers the labor share. As in our model, this reflects an increase in output that more than offset the rise in labor income. Moreover, as in our findings, the negative effect on the labor share has been growing over time and was more pronounced since 2000.[2]

Finally, our model has important implications for the volatility of the vacancy-unemployment ratio relative to the volatility of real wages, a statistical moment that has received a lot of attention in the literature. In our framework, the threat of automation endogeneously dampens wage fluctuations relative to unemployment and vacancies, implicitly working as a form of real wage rigidity. This mechanism is quantitatively important. In our estimated model, the volatility of the vacancy-unemployment ratio (i.e., the v-u ratio), which is a measure of labor market tightness, is about 40 times that of the real wage rate, much larger than that obtained from a counterfactual model with no endogenous fluctuations in automation (with a volatility ratio of about 9). The procyclical threat of automation dampens wage increases. At the same time, the increased adoption of robots raises aggregate productivity. Through these channels, our model is able to generate a large volatility of the labor market tightness relative to that of the real wage rate.

Our macro approach on the role automation complements recent analyses using disaggregated data. For instance, Arnoud (2018) examines the threat of automation on wages using data from the U.S. Current Population Survey in 2013 and an index of automatability developed in the literature (Frey and Osborne, 2017). Looking at wages by occupations and individual characteristics, he finds a negative effect. Our model's prediction that the threat of automation mutes wages is consistent with this micro-level evidence.

Dinlersoz and Wolf (2018) use data from three U.S. Census Bureau's Survey of Manufacturing Technology conducted in the late 1980s and early 1990s. One appealing feature of these surveys for studying automation is that they were specifically designed to look at technologies that facilitated the substitution of capital for labor. In line with our findings,

[2]Because their sample ends in 2007, Author and Salomons are silent on the impact of automation on the labor share following the Great Recession.

these authors document that firms that made greater use of automation displayed higher productivity and had a lower labor share.

The impact of a greater use of robots in production has also been examined by Acemoglu and Restrepo (2017), using U.S. county-level employment data on robots adoption, finding a negative impact on employment and wages. Using a panel of industries in 17 countries from 1993 to 2007, Graetz and Michaels (2018) estimate that increases in the use of robots boosts labor productivity growth and wages, with little effect on aggregate employment (though it tends to reduce employment of low-skilled workers).

The empirical literature using disaggregated data is well suited to look at the impact of automation on different types of industries, jobs, and tasks. However, one drawback is the difficulty to aggregate this micro evidence into a macroeconomic impact. Our DSGE model instead embeds all the general equilibrium effects, but does not directly speak to heterogeneous effects across different types of jobs or sectors.

Karabarbounis and Neiman (2013) document a general decline in the labor share across 59 countries between 1975 and 2012. They introduce a macro model in which a secular fall in the relative price of investment give firms an incentive to substitute capital for labor. In their benchmark calibration, which includes an elasticity of substitution between capital and labor greater than one, they find that automation accounts for roughly half of the decline in the labor share in their sample. Our approach complements this work by focusing on the relation between automation and wage bargaining and by focusing on business cycle fluctuations.

There are other factors that may have contributed to the decline in the labor share since the early 2000s. In addition to automation, Elsby et al. (2013) also empirically examine the role of offshoring and find that industries facing greater import competition, particularly from China, experienced a larger decline in the labor share. Autor et al. (2019) use micro, firm-level, data form U.S. censuses since 1982 to calculate industry concentration and examine the role of the rise in superstar firms for the decline in the labor share. They finding a larger fall in the labor share in industries that experienced greater concentration over time. Relatedly, Krueger (2018) emphasizes the growing importance of firms' monopsony power in the labor market (in the forms of implicit collusions, non-compete restrictions in labor contracts, and outsourcing) as a factor that diminishes workers' bargaining power. Our model highlights a different mechanism: even if the worker's bargaining weight is held constant, a higher threat of automation can reduce wages because it raises the firms' reservation value. Thus, our work emphasizes the importance of regarding automation and bargaining power as being jointly determined.

The reminder of the paper is organized as follows. We first introduce the model in the next section and then present our empirical strategy. We next discuss the business cycle contributions of the different shocks in our model before examining the economy's responses to these shocks and the role that automation plays in the transmission. Using outputs from our estimation strategy, we quantify the effect of the threat of automation on the labor share, wages, and productivity. The last section concludes.

II. The model with labor market frictions and automation

This section presents a DSGE model that generalizes the standard DMP model to incorporate endogenous decisions of automation.

To keep automation decisions tractable, we impose some assumptions on the timing of events. In the beginning of period t, a job separation shock δ_t realizes. Workers who lose their jobs adds to the stock of unemployment from the previous period, forming the pool of job seekers u_t. Firms post vacancies v_t at a fixed cost κ. The stock of vacancies v_t includes the unfilled vacancies that were not automated at the end of period $t-1$, the jobs separated in the beginning of period t, and new vacancies created in the beginning of period t. Creating a new vacancy incurs a fixed cost, which is drawn from an i.i.d. distribution $G(\cdot)$ as in Leduc and Liu (2019). In the labor market, a matching technology transforms job seekers and vacancies into an employment relation, with a wage rate determined through Nash bargaining between the employer and the job seeker. Once an employment relation is formed, production takes place, and the firm receives the employment value. An unfilled vacancy can be either carried forward to the next period or automated at a fixed cost. Similar to the vacancy creation cost, the automation cost x is drawn from an i.i.d. distribution $F(x)$. If a firm draws an automation cost that is below a threshold value x_t^*, then the firm adopts a robot and closes the job opening. In that case, the firm obtains the automation value. Otherwise, the vacancy remains open and the firm receives the continuation value of the vacancy. Newly adopted robots add to the stock of automation, which becomes obsolete over time at a constant rate ρ^o. Final goods output is the sum of the goods produced by workers and by robots. The final good is used for household consumption and also for paying the costs of vacancy posting, new vacancy creation, and robot adoption.

II.1. **The Labor Market.**

In the beginning of period t, there are N_{t-1} existing job matches. A job separation shock displaces a fraction δ_t of those matches, so that the measure of unemployed job seekers is given by

$$u_t = 1 - (1-\delta_t)N_{t-1}, \tag{1}$$

where we have assumed full labor force participation and normalized the size of the labor force to one.

The job separation rate shock δ_t follows the stationary stochastic process

$$\ln \delta_t = (1 - \rho_\delta) \ln \bar{\delta} + \rho_\delta \ln \delta_{t-1} + \varepsilon_{\delta t}, \tag{2}$$

where ρ_δ is the persistence parameter and the term $\varepsilon_{\delta t}$ is an i.i.d. normal process with a mean of zero and a standard deviation of σ_δ. The term $\bar{\delta}$ denotes the steady-state rate of job separation.

The stock of vacancies v_t in the beginning of period t consists of the vacancies in period $t-1$ that were not filled with workers and not automated, plus the displaced job positions and newly created vacancies. The law of motion for vacancies is given by

$$v_t = (1 - q^v_{t-1})(1 - q^a_{t-1}))v_{t-1} + \delta_t N_{t-1} + \eta_t, \tag{3}$$

where q^v_{t-1} denotes the job filling rate in period $t-1$, q^a_{t-1} denotes the automation rate in period $t-1$, and η_t denotes the newly created vacancies (i.e., entry).

In the labor markert, new job matches are formed between job seekers and open vacancies based on the matching function

$$m_t = \mu u_t^\alpha v_t^{1-\alpha}, \tag{4}$$

where m_t denotes the number of job matches (or hiring) and $\alpha \in (0, 1)$ is the elasticity of job matches with respect to the number of job seekers.

The flow of new job matches adds to the employment pool and job separations subtract from it. Aggregate employment evolves according to the law of motion

$$N_t = (1 - \delta_t)N_{t-1} + m_t. \tag{5}$$

At the end of period t, the searching workers who failed to find a job match remain unemployed. Thus, unemployment is given by

$$U_t = u_t - m_t = 1 - N_t. \tag{6}$$

For convenience, we define the job finding probability q^u_t as

$$q^u_t = \frac{m_t}{u_t}. \tag{7}$$

Similar, we define the job filling probability q^v_t as

$$q^v_t = \frac{m_t}{v_t}. \tag{8}$$

II.2. **The firms.** If a firm successfully hires a worker, then it can produce Z_t units of intermediate goods. The technology shock Z_t follows the stochastic process

$$\ln Z_t = (1-\rho_z)\ln \bar{Z} + \rho_z \ln Z_{t-1} + \varepsilon_{zt}. \tag{9}$$

The parameter $\rho_z \in (-1,1)$ measures the persistence of the technology shock. The term ε_{zt} is an i.i.d. normal process with a zero mean and a finite variance of σ_z^2. The term $\bar{Z}$ is the steady-state level of the technology shock.[3]

The value of employment satisfies the Bellman equation

$$J_t^e = Z_t - w_t + \mathbb{E}_t D_{t,t+1}\left\{(1-\delta_{t+1})J_{t+1}^e + \delta_{t+1}J_{t+1}^v\right\}, \tag{10}$$

where $D_{t,t+1}$ is a stochastic discount factor of the households. Hiring a worker generates a flow profit $Z_t - w_t$ in the current period. If the job is separated in the next period (with probability δ_{t+1}), then the firm receives the vacancy value J_{t+1}^v. Otherwise, the firm receives the continuation value of employment.

Following Leduc and Liu (2019), we assume that creating a new vacancy incurs an entry cost e in units of consumption goods. The entry cost is drawn from an i.i.d. distribution $F(e)$. A new vacancy is created if and only if the net value of entry is non-negative. The benefit of creating a new vacancy is the vacancy value J_t^v. Thus, the number of new vacancies η_t is given by the cumulative density of the entry costs evaluated at J_t^v. That is,

$$\eta_t = F(J_t^v). \tag{11}$$

Posting a vacancy incurs a per-period fixed cost κ (in units of final consumption goods). If the vacancy is filled (with probability q_t^v), the firm obtains the employment value J_t^e. If the vacancy is not filled, then the firm can choose to automate the position and close the vacancy (with probability q_t^a), in which case the firm obtains the automation value J_t^a. If the firm does not automate the unfilled position, then it receives the continuation value of the vacancy. Thus, the vacancy value satisfies the Bellman equation

$$J_t^v = -\kappa + q_t^v J_t^e + (1-q_t^v)q_t^a J_t^a + (1-q_t^v)(1-q_t^a)\mathbb{E}_t D_{t,t+1}J_{t+1}^v. \tag{12}$$

The flow of automated job positions adds to the stock of automation, which becomes obsolete at the rate $\rho^o \in [0,1]$ in each period. Thus, the automation stock A_t evolves according to the law of motion

$$A_t = (1-\rho^o)A_{t-1} + q_t^a(1-q_t^v)v_t, \tag{13}$$

where $q_t^a(1-q_t^v)v_t$ is the number of unfilled job positions that are newly automated in period t.

[3]The model can easily be extended to allow for trend growth. We do not present that version of the model to simplify presentation.

Once adopted, a robot produced $Z_t\zeta_t$ units of output, where ζ_t denotes an equipment-specific technology shock, which follows a stochastic process that is independent of the neutral technology shock Z_t. In particular, ζ_t follows the stationary process

$$\ln \zeta_t = (1-\rho_\zeta)\ln \bar{\zeta} + \rho_\zeta \ln \zeta_{t-1} + \varepsilon_{\zeta t}. \tag{14}$$

The parameter $\rho_\zeta \in (-1,1)$ measures the persistence of the automation-specific technology shock. The term $\varepsilon_{\zeta t}$ is an i.i.d. normal process with a zero mean and a finite variance of σ_ζ^2. The term $\bar{\zeta}$ is the steady-state level of the automation-specific technology shock.

Operating the robot incurs a flow fixed cost of κ_a. The value of automation satisfies the Bellman equation

$$J_t^a = Z_t\zeta_t - \kappa_a + (1-\rho^o)\mathbb{E}_t D_{t,t+1} J_{t+1}^a, \tag{15}$$

where the term κ_a captures the costs of energy, facilities, and space for automated production.

Automating a vacancy requires a fixed cost x in units of consumption goods. The fixed cost is drawn from the i.i.d. distribution $G(x)$. A firm chooses to adopt a robot if and only if the cost of automation is less than the benefit. For any given benefit of automation, there exists a threshold value x_t^* in the support of the distribution $G(x)$, such that automation occurs if and only if $x \leq x_t^*$. The threshold value x_t^* depends on the value of automation J_t^a relative to the continuation value of a vacancy. In particular, the threshold for automation decision is given by

$$x_t^* = J_t^a - \mathbb{E}_t D_{t,t+1} J_{t+1}^v. \tag{16}$$

Thus, the probability of automation is the cumulative density of the automation costs evaluated at x_t^*. More formally, the automation probability is determined by

$$q_t^a = G(x_t^*). \tag{17}$$

II.3. **The representative household.** The representative household has the utility function

$$\mathbb{E}\sum_{t=0}^{\infty} \beta^t \Theta_t \left(\ln C_t - \chi N_t\right), \tag{18}$$

where $\mathbb{E}[\cdot]$ is an expectation operator, C_t denotes consumption, and N_t denotes the fraction of household members who are employed. The parameter $\beta \in (0,1)$ denotes the subjective discount factor, and the term Θ_t denotes an exogenous shifter to the subjective discount factor.

The discount factor shock $\theta_t \equiv \frac{\Theta_t}{\Theta_{t-1}}$ follows the stationary stochastic process

$$\ln \theta_t = \rho_\theta \ln \theta_{t-1} + \varepsilon_{\theta t}. \tag{19}$$

In this shock process, ρ_θ is the persistence parameter and the term $\varepsilon_{\theta t}$ is an i.i.d. normal process with a mean of zero and a standard deviation of σ_θ. Here, we have implicitly assumed that the mean value of θ is one.

The representative household chooses consumption C_t and savings B_t to maximize the utility function in (18) subject to the sequence of budget constraints

$$C_t + \frac{B_t}{r_t} = B_{t-1} + w_t N_t + \phi(1 - N_t) + d_t - T_t, \quad \forall t \geq 0, \tag{20}$$

where r_t denotes the gross real interest rate, w_t denotes the real wage rate, d_t denotes the household's share of firm profits, and T_t denotes lump-sum taxes. The parameter ϕ measures the flow benefits of unemployment.

Denote by $V_t(B_{t-1}, N_{t-1})$ the value function for the representative household. The household's optimizing problem can be written in the recursive form

$$V_t(B_{t-1}, N_{t-1}) \equiv \max \ \ln C_t - \chi N_t + \beta \mathbb{E}_t \theta_{t+1} V_{t+1}(B_t, N_t), \tag{21}$$

subject to the budget constraint (20) and the employment law of motion (5), the latter of which can be written as

$$N_t = (1 - \delta_t) N_{t-1} + q^u u_t, \tag{22}$$

where we have used the definition of the job finding probability $q_t^u = \frac{m_t}{u_t}$, with the measure of job seekers u_t given by Eq. (1). In the optimizing decisions, the household takes the economy-wide job finding rate q_t^u as given.

Define the employment surplus (i.e., the value of employment relative to unemployment) as $S_t^H \equiv \frac{1}{\Lambda_t} \frac{\partial V_t(B_{t-1}, N_{t-1})}{\partial N_t}$, where Λ_t denotes the Lagrangian multiplier for the budget constraint (20). We show in the Appendix that the employment surplus satisfies the Bellmand equation

$$S_t^H = w_t - \phi - \frac{\chi}{\Lambda_t} + \mathbb{E}_t D_{t,t+1}(1 - q_{t+1}^u)(1 - \delta_{t+1}) S_{t+1}^H, \tag{23}$$

where $D_{t,t+1} \equiv \frac{\beta \theta_{t+1} \Lambda_{t+1}}{\Lambda_t}$ is the stochastic discount factor, which applies to both the household's intertemporal optimization and the firms' decisions.

The employment surplus has a straightforward economic interpretation. If the household adds a new worker in period t, then the current-period gain would be wage income net of the opportunity costs of working, including unemployment benefit and the disutility of working. The household also enjoys the continuation value of employment if the employment relation continues. Having an extra worker today adds to the employment pool tomorrow (if the employment relation survives job separation); however, adding a worker today would also reduce the pool of searching workers tomorrow, a fraction q_{t+1}^u of whom would be able to find jobs. Thus, the marginal effect of adding a new worker in period t on employment in period $t+1$ is given by $(1-q^u_{t+1})(1-\delta_{t+1})$, resulting in the effective continuation value of employment shown in the last term of Eq. (23).

We also show in the appendix that the household's optimizing consumption-savings decision implies the intertemporal Euler equation

$$1 = \mathbb{E}_t D_{t,t+1} r_t. \tag{24}$$

II.4. **The Nash bargaining wage.** When a job match is formed, the wage rate is determined through Nash bargaining. The bargaining wage optimally splits of the joint surplus of a job match between the worker and the firm. The worker's employment surplus is given by S^H_t in Eq. (23). The firm's surplus is given by $J^e_t - J^v_t$. The possibility of automation affects the value of a vacancy, and thus indirectly affects the firm's reservation value and their bargaining decisions.

The Nash bargaining problem is given by

$$\max_{w_t} \quad \left(S^H_t\right)^b \left(J^e_t - J^v_t\right)^{1-b}, \tag{25}$$

where $b \in (0,1)$ represents the bargaining weight for workers.

Define the total surplus as

$$S_t \equiv J^e_t - J^v_t + S^H_t. \tag{26}$$

Then the bargaining solution is given by

$$J^e_t - J^v_t = (1-b)S_t, \quad S^H_t = bS_t. \tag{27}$$

The bargaining outcome implies that the firm's surplus is a constant fraction $1-b$ of the total surplus S_t and the household's surplus is a fraction b of the total surplus.

The bargaining solution (27) and the expression for household surplus in equation (23) together imply that the Nash bargaining wage w^N_t satisfies the Bellman equation

$$\begin{aligned} \frac{b}{1-b}(J^e_t - J^v_t) &= w^N_t - \phi - \frac{\chi}{\Lambda_t} \\ &+ \mathbb{E}_t D_{t,t+1}(1-q^u_{t+1})(1-\delta_{t+1})\frac{b}{1-b}(J^e_{t+1} - J^v_{t+1}). \end{aligned} \tag{28}$$

II.5. **Government policy.** The government finances unemployment benefit payments ϕ for unemployed workers through lump-sum taxes. We assume that the government balances the budget in each period so that

$$\phi(1-N_t) = T_t. \tag{29}$$

II.6. **Search equilibrium.** In a search equilibrium, the markets for bonds and goods all clear. Since the aggregate bond supply is zero, the bond market-clearing condition implies that

$$B_t = 0. \tag{30}$$

Goods market clearing requires that consumption spending, automation costs, and vacancy creation costs add up to aggregate production. This requirement yields the aggregate resource constraint

$$C_t + \kappa v_t + (1 - q_t^v) v_t \int_0^{x_t^*} x dG(x) + \int_0^{J_t^v} e dF(e) = Y_t, \tag{31}$$

where Y_t denotes aggregate output.

Aggregate output is sum of goods produced by workers and by robots. Specifically, it is given by

$$Y_t = Z_t N_t + Z_t \zeta_t A_t. \tag{32}$$

III. Empirical Strategies

We solve the model by log-linearizing the equilibrium conditions around the deterministic steady state.[4] We calibrate a subset of the parameters to match steady-state observations and the empirical literature. We estimate the remaining structural parameters and the shock processes to fit U.S. time-series data.

We focus on the parameterized distribution functions

$$F(e) = \left(\frac{e}{\bar{e}}\right)^{\eta_v}, \quad G(x) = \left(\frac{x}{\bar{x}}\right)^{\eta_a}, \tag{33}$$

where $\bar{e} > 0$ and $\bar{x} > 0$ are the scale parameters and $\eta_v > 0$ and $\eta_a > 0$ are the shape parameters of the distribution functions. We set $\eta_v = 1$ and $\eta_a = 1$, so that both the vacancy creation cost and the automation cost follow a uniform distribution.[5] We estimate the scale parameters $\bar{e}$ and $\bar{x}$ and the shock processes by fitting the model to U.S. time series data.

[4]Details of the equilibrium conditions, the steady state, and the log-linearized system are presented in the appendix.

[5]Our assumption of the uniform distribution for the vacancy creation cost is in line with Fujita and Ramey (2007). We have estimated a version of the model in which we include the parameter η_a in the set of parameters to be estimated. We obtain a posterior estimate of $\eta_a = 1.007$ and very similar estimates for the other parameters. For simplicity and ofr obtaining closed-form solution for the steady state equilibrium we assume that $\eta_a = 1$ in our benchmark model.

III.1. **Steady state equilibrium and parameter calibration.** Table 1 shows the calibrated parameter values. We consider a quarterly model. We set $\beta = 0.99$, so that the model implies a annualized real interest rate of about 4 percent in the steady state. We set $\alpha = 0.5$ following the literature (Blanchard and Galí, 2010; Gertler and Trigari, 2009). In line with Hall and Milgrom (2008), we set $b = 0.5$ and $\phi = 0.25$. Based on the data from JOLTS, we calibrate the steady-state job separation rate to $\bar{\delta} = 0.10$ at the quarterly frequency. We set $\rho^o = 0.03$, so that robots depreciate at an average annual rate of 12 percent. We normalize the level of labor productivity to $\bar{Z} = 1$ and automation-specific productivity to $\bar{\zeta} = 1$.

We target a steady-state unemployment rate of $U = 0.058$, corresponding to the average unemployment rate in the period from 1951 to 2018. The steady-state employment is given by $N = 1 - U$, hiring rate by $m = \bar{\delta}N$, and the number of job seekers by $u = 1 - (1 - \bar{\delta})N$. The job finding rate is $q^u = \frac{m}{u}$. We target a steady-state job filling rate q^v of 0.71 per quarter, in line with the calibration of den Haan et al. (2000). The implied stock of vacancies is $v = \frac{m}{q^v}$. The scale of the matching efficiency is then given by $\mu = \frac{m}{u^\alpha v^{1-\alpha}} = 0.6629$. We set the flow cost of operating robots to $\kappa_a = 0.98$. Given the average productivities $\bar{Z} = \bar{\zeta} = 1$, this implies a quarterly profit of 2 percent of the revenue by using a robot for production. The steady-state automation value J^a can then be solved from the Bellman equation (15).

Conditional on J^a and the estimated values of $\bar{e}$ and $\bar{x}$ (see below for estimation details), we use the vacancy creation condition (11), the automation adoption condition (16), and law of motion for vacancies (3) to obtain the steady-state probability of automation, which is given by

$$q^a = \frac{J^a}{\bar{x} + \beta\bar{e}(1 - q^v)v}.$$

Given q^a and v, the law of motion for vacancies implies that the flow of new vacancies is given by $\eta = q^a(1 - q^v)v$. The vacancy value is then given by $J^v = \bar{e}\eta$. The stock of automation A can be solved from the law of motion (13), which reduces to $\rho^o A = q^a(1 - q^v)v = \eta$ in the steady state. Thus, in the steady state, the newly created vacancies equal the flow of automated jobs that become obsolete. The law of motion for employment implies that, in the steady state, the flow of hiring equals the flow of separated employment relations.

With A and N solved, we obtain the aggregate output $Y = \bar{Z}(N + \bar{\zeta}A)$. We calibrate the vacancy posting cost to $\kappa = 0.0884$, so that the steady-state vacancy posting cost is one percent of aggregate output (i.e., $\kappa v = 0.01Y$).

Given J^v and J^a, we obtain the cutoff point for robot adoption $x^* = J^a - \beta J^v$. The match value J^e can be solved from the Bellman equation for vacancies (12), and the equilibrium real wage rate can be obtained from the Bellman equation for employment (10). Steady-state consumption is solved from the resource constraint (31). We then infer the value of $\chi = 0.7267$ from the expression for bargaining surplus in Eq. (28).

III.2. **Estimation.** We now describe our estimation approach.

III.2.1. *Data and measurement.* We fit the DSGE model to four quarterly U.S. time series: the unemployment rate, the job vacancy rate, the growth rate of average labor productivity in the non-farm business sector, and the growth rate of the real wage rate. The sample covers the range from 1951:Q1 to 2018:Q4.

The unemployment rate in the data (denoted by U_t^{data}) corresponds to the end-of-period unemployment rate in the model U_t. We demean the unemployment rate data (in log units) and relate it to our model variable according to

$$\ln(U_t^{data}) - \ln(\bar{U}^{data}) = \hat{U}_t, \tag{34}$$

where $\bar{U}^{data}$ denotes the sample average of the unemployment rate in the data and $\hat{U}_t$ denotes the log-deviations of the unemployment rate in the model from its steady-state value.

Similarly, we use demeaned vacancy rate data (also in log units) and relate it to the model variable according to

$$\ln(v_t^{data}) - \ln(\bar{v}^{data}) = \hat{v}_t, \tag{35}$$

where $\bar{v}^{data}$ denotes the sample average of the vacancy rate data and $\hat{v}_t$ denotes the log-deviations of the vacancy rate in the model from its steady-state value. Our vacancy series for the periods prior to 2001 is the vacancy rate constructed by Barnichon (2010) based on the Help Wanted Index. For the periods after 2001, we use the JOLTS vacancy rate.

In the data, we measure labor productivity by real output per person in the non-farm business sector. We use the demeaned quarterly log-growth rate of labor productivity (denoted by $\Delta \ln p_t^{data}$) and relate it to our model variable according to

$$\Delta \ln(p_t^{data}) - \Delta \ln(p^{data}) = \hat{Y}_t - \hat{N}_t - (\hat{Y_{t-1}} - \hat{N_{t-1}}), \tag{36}$$

where $\Delta \ln(p^{data})$ denotes the sample average of productivity growth, and $\hat{Y}_t$ and $\hat{N}_t$ denote the log-deviations of aggregate output and employment from their steady-state levels in our model.

We measure the real wage rate in the data by the real compensation per worker in the non-farm business sector. We relate the observed real wage growth (denoted by $\Delta \ln(w_t^{data})$) to the model variables by the measurement equation

$$\Delta \ln(w_t^{data}) - \Delta \ln(w^{data}) = \hat{w}_t - \hat{w}_{t-1}, \tag{37}$$

where $\Delta \ln(w^{data})$ denotes the sample average of wage growth in the data and $\hat{w}_t$ denotes the log-deviations of real wages from its steady-state level in the model.

III.2.2. *Prior distributions and posterior estimates.* The prior and posterior distributions of the estimated parameters from our benchmark model are displayed in Table 2.

The priors for the structural parameters $\bar{e}$ and $\bar{x}$ are drawn from the gamma distribution. We assume that the prior mean of each of these three parameters is 5, with a standard deviation of 1. The priors of the persistence parameter of each shock follow the beta distribution with a mean of 0.8 and a standard deviation of 0.1. The priors of the volatility parameter of each shock follow an inverse gamma distribution with a prior mean of 0.01 and a standard deviation of 0.1.

The posterior estimates and the 90 percent probability intervals for the posterior distributions are displayed in the last three columns of Table 2. The posterior mean estimate of the vacancy creation cost parameter is $\bar{e} = 9.5693$. The posterior mean estimates of the automation cost parameter is $\bar{x} = 2.4349$. These parameters imply a steady-state automation probability of about 18 percent per quarter, which lies in the range estimated in the empirical literature (Frey and Osborne, 2015; Nedelkoska and Quintini, 2018). The 90-percent probability intervals indicate that the data are informative about these structural parameters.

The posterior estimation suggests that the shocks to both the neutral technology and the discount factor are highly persistent, whereas the automation-specific shock is less persistent but more volatile. The 90-percent probability intervals suggest that the data are also informative for these shock processes.

IV. Economic implications

We now examine the model's transmission mechanism and its quantitative performance for explaining the labor market dynamics.

IV.1. **The model's transmission mechanism.** The equilibrium dynamics in our model are driven by both the exogenous shocks and the model's internal propagation mechanism. To help understand the contributions of the shocks and the model's mechanism, we examine forecast error variance decompositions and impulse response functions.

IV.1.1. *Forecast error variance decompositions.* Table 3 displays the unconditional forecast error variance decompositions for the four observable labor market variables used for our estimation.[6]

The variance decompositions suggest that the fluctuations of unemployment and vacancies are mostly driven by the neutral technology shock and the discount factor shock. In line with

[6]We have also computed the conditional forecast error variance decompositions with forecasting horizons between 4 quarters and 16 quarters and found that they deliver the same message as the unconditional variance decomposition.

the literature (Hall, 2017; Leduc and Liu, 2019), the discount factor shock accounts for over half of the variances of unemployment and vacancies. In our model, the procyclical threat of automation dampens real wage adjustments and thus amplifies the impact of technology shocks on labor market variables. The amplification mechanism is quantitatively important. As shown in Table 3, the neutral technology shock accounts for over 40 percent of the fluctuations in unemployment and vacancies. The job separation shock is not important for these labor market variables, consistent with Shimer (2005).

The automation-specific shock does not directly contribute much to the fluctuations in unemployment and vacancies; instead, the threat of automation works to amplify the effects of the other shocks. This amplification mechanism works through endogenous variations in the probability of automation, driven by other shocks such as the neutral technology shock and the discount factor shock. As shown in the table, these two shocks are both important, and they together explain over 95 percent of the fluctuations in the automation probability.

In our model, goods produced by workers and by robots are perfect substitutes. Thus, the actual adoption of robots has a direct job-displacement effect. However, automation also boosts employment through two channels. First, the option of automating an unfilled job position raises the present value of a vacancy, encouraging firms to create more vacancies. Second, the threat of automation strengthens the firm's bargaining power in wage negotiations because it raises the firm's outside option, leading to muted wage increases in a business cycle boom. This endogenous real wage rigidity helps amplify the responses of unemployment and vacancies to technology shocks and discount factor shocks. The general equilibrium job-creating effects of automation counteracts the direct job-displacing effects, and the net effect on employment depends on parameter values. With our estimated parameters, we show that automation amplifies employment fluctuations. In the next section, we will elaborate the specific channels for the general equilibrium effects of automation with the help of impulse responses.

While the threat of automation dampens wage adjustments, the actual adoption of robots raises labor productivity. Through these channels, the automation-specific shock plays a quantitatively important role in driving fluctuations of the growth rates of both labor productivity and real wages. The shock accounts for about 33 percent of the variance of productivity growth and 49 percent of that of real wage growth.

Perhaps not surprisingly, the neutral technology shock is important for labor productivity fluctuations, explaining about half of its variance.[7] The shock also accounts for about a

[7] In the standard DMP version of our model without automation, labor productivity fluctuations would be entirely driven by the neutral technology shock.

quarter of the variance of real wage growth. The discount factor shock explains the remaining fluctuations (about 20 percent) in labor productivity and real wage growth.

IV.1.2. *Impulse responses.* To understand the transmission mechanism of the model, we examine the impulse responses of several key labor market variables following each shock. We also highlight the role of automation in the model's transmission mechanism by examining some counterfactual models without automation.

Figures 1 shows the impulse responses to a positive neutral technology shock in the benchmark model. The shock leads to persistent declines in unemployment and persistent increases in vacancies and hiring. The neutral technology shock increases the productivity of all jobs (operated by a worker or a robot), so that the net value of automation rises, leading to an increased probability of robot adoption. Since an unfilled job position can be automated more easily, the vacancy value rises, encouraging firms to create more vacancies. The increase in vacancy value also strengthens the firm's bargaining power in wage negotiations, dampening the responses of real wages. With more jobs automated, labor productivity increases persistently, reinforcing the initial postive impact of the shock. The increase in labor productivity, coupled with muted wage responses, implies persistent declines in the labor income share.

Automation amplifies not just the neutral technology shock, but also the discount factor shock, as shown in Figure 2. A positive discount factor shock generates a persistent boom in employment, vacancies, and hiring by raising the present values of a job match, an open vacancy, and a worker's employment surplus. The shock also raises the present value of automation, and therefore leading to an increase in the probability of robot adoption. The increased probability of automation raises the present value of a vacancy, stimulating the incentive to create new vacancies. With more vacancies created, the job finding rate increases, amplifying the employment boom. Following a positive discount factor shock, the increased threat of automation leads to a modest short-run decline in the real wage rate. Although the shock originates from the demand side, it has a positive impact on labor productivity because of the increased adoption of automation. With muted wage responses, the increases in labor productivity leads to a decline in the labor income share.

The job separation shock raises both unemployment and vacancies and mechanically boosts hiring through the matching function, as shown in Figure 3. This finding is consistent with Shimer (2005), who argues that the counterfactual implication of the job separation shock on the correlation between unemployment and vacancies renders the shock unimportant for explaining the observed labor market dynamics. The shock reduces the automation probability because it increases the vacancy value without affecting the automation value. Labor productivity increases slightly because the decline in employment outpaces the decline in aggregate output, a part of which is produced with robots. The shock also leads small declines in real wages and the labor income share.

Figure 4 shows the impulse responses to a positive automation-specific shock. The shock directly raises the value of automation. The increased probability of automation raises the vacancy value and boosts the incentive for vacancy creation. With more job openings, the job finding rate increases, raising hiring and reducing unemployment. Since a greater fraction of output is produced with robots, labor productivity improves. The increased threat of automation weakens the worker's bargaining power, leading to a decline in the real wage rate. The improvement in labor productivity and the reduction in the real wage rate lead to a persistent decline in the labor income share.

IV.2. **The role of automation in the propagation mechanism.** To isolate the role of endogenous automation decisions in driving the labor market dynamics, we consider a counterfactual specification of "no automation," which is a version of our benchmark model with all automation-related variables held constant at their steady-state levels and with no automation-specific shocks. To highlight the effect of the threat of automation on the worker's bargaining power in wage negotiations, we also compare our benchmark model's impulse responses to a version of the "no automation" case, in which we also reduce the bargaining weight for workers by a half (i.e., setting $b = 0.25$).

Figure 6 displays the impulse responses to a discount factor shock in the three models: the benchmark model (the black solid lines), the counterfactual with no automation (the blue dashed lines), and the counterfactual with no automation and a lower bargaining power for workers (the red dashed and dotted lines).[8] In the counterfactual with no automation, the responses of unemployment, vacancies, and hiring to the shock are much more muted than in the benchmark model. Without automation, lowering the workers' bargaining power helps amplify the labor market responses, but the magnitude of amplification is dwarfed by that of the automation mechanism. In the counterfactuals without automation, real wages rise slightly following the discount factor shock. In contrast, the threat of automation in our benchmark model leads to modest short-run declines in the real wage rate. Without automation, labor productivity is solely driven by the neutral technology shock, and it does not respond to the discount factor shock. Thus, without automation, the labor income share rises slightly, reflecting the increases in the real wage rate. The no-automation model generates procyclical labor income share even if the real wage responses are mechanically muted by reducing the workers' bargaining power. With automation, our benchmark model

[8]We report the impulse responses to the neutral technology shock and the job separation shock for the counterfactuals in the appendix. Since automation is turned off in both counterfactual specifications, the automation-specific shock is irrelevant for studying these counterfactuals.

predicts that the discount factor shock raises the labor productivity persistently since the shock raises the value of automation and leads to more robot adoptions. The increase in labor productivity and the modest declines in the real wage rate imply a countercyclical labor income share.

IV.2.1. *Automation threat and labor market dynamics.* Our model predicts that the threat of automation dampens wage adjustments and amplifies labor market fluctuations. By raising labor productivity while muting wage responses, automation also implies a countercyclical labor income share. Is the automation mechanism quantitatively important? To examine the empirical importance of the automation mechanism, we compare our model's predictions for the unemployment rate and the labor income share with those from counterfactuals without the automation channel.

Our estimated model implies that the probability of automation is procylical, rising in business cycle booms and falling in recessions. Figure 7 shows the smoothed series of the automation probability obtained from our estimated model, and confirms the procyclical behavior of the automation probability.

Since the option of automation raises the value of a job vacancy, it encourages job creation, which more than offset the direct job-displacing effect of automation, amplifying unemployment fluctuations. This is illustrated by Figure 8, which shows the smoothed series of unemployment from our estimated benchmark model (the blue solid line) and that from a counterfactual model with no automation (the red dashed line). The smoothed series from our benchmark model are identical to the actual series in the data because we estimate our model to fit those series. The "no automation" counterfactual is a version of our benchmark model with all the automation-related variables held constant at their steady-state values and without automation-specific shocks. The gap between the blue line and the red line shows the contribution of the automation mechanism and the automation shock to unemployment fluctuations. Evidently, procyclical automation threat is crucial for explaining the observed unemployment dynamics. For example, after the 2008-09 Great Recession, the model without automation predicts incorrectly that the unemployment rate would have reached a peak of 6.5 percent by the end of 2011, while the actual peak was much higher at more than 10 percent, and it occurred at least two years earlier. The counterfactual model without automation also missed the magnitude and persistence of the declines in the unemployment rate in the recent few years: it predicts that the unemployment rate should stay at about 5.5 percent by the end of 2018, whereas in the data, the unemployment rate was under 4 percent. Thus, the automation mechanism significantly amplified procyclical fluctuations in employment.

Figure 9 shows the labor income share implied by our estimated model (the blue solid line) and that from a counterfactual model without automation (the red dashed line). Since we estimate the benchmark model to fit both labor productivity growth and real wage growth, we can recover the labor income share from the smoothed series of real wages and labor productivity. We do the same for the counterfactual model with no automation, but otherwise with the same parameters and shock processes. The figure shows that automation has played a quantitatively important role in explaining the observed labor share fluctuations. In the counterfactual model without automation, the predicted labor income share (the red line) declined steadily throughout the period from 1980 to 2018. The actual labor share (the blue line) has also declined since 1980, and the decline accelerated sharply since the early 2000s. The counterfactual path of the labor share has trended down together with the actual labor share until 2009 (although it misses the short-run fluctuations). After 2009, the two paths diverged: while the actual labor share fell sharply, the counterfactual model predicts that the labor share in 2018 should have stayed at roughly the pre-recession level.

The divergence between the two paths after 2009 highlights the importance of the automation channel in explaining the observed labor share dynamics. As illustrated by Figure 7, the post-2009 period experienced high automation probability, The increased automation probability raises the firm's reservation value in wage bargaining and dampens wage increases. Indeed, as shown in Figure 10, automation has contributed to lowering real wages since the early 2000s, with the wage-depressing effects particularly pronounced after the Great Recession. At the same time, the increased robot adoption has boosted labor productivity since the early 2000s, as shown in Figure 11. Thus, with the automation channel turned on, our benchmark model predicts a sharp decline in the labor share.

Of course, the actual labor productivity growth in the post-recession period has been significantly slower than that observed in the late 1990s and 2000s (Fernald, 2015). From the lens of our model, this slowdown in labor productivity growth is driven by persistent negative shocks to the neutral technology, partly offset by positive automation shocks. Had there been no increases in automation, labor productivity growth would have been even slower since the mid-2000s.

V. Conclusion

We have studied how automation interacts with labor market dynamics in a DMP framework that incorporates endogenous automation decisions. The option to automate a job position induces a job-ceation incentive, which offsets the direct job-displacing effects of automation, and can thus potentially amply unemployment fluctuations. Procyclical threat of automation raises the firm's reservation value in wage bargaining, dampening increases in real wages in a business cycle boom, also amplying fluctuations in unemployment and vacancies. Furthemore, actual adoptions of robots raise labor productivity. With muted wage increases and amplified labor productivity, automation leads to declines in the labor share in good economic times.

To assess the quantitative importance of the automation channel, we estimate our DMP model with automation to fit the U.S. time series data. We find that automation contributed significantly to the observed sluggish wage growth and the rapid declines in the labor share during the post-Great Recession period.

Of course, many other factors may have contributed to the wage and labor share dynamics since the early 2000s. Examples include declines in workers' bargaining power, increases in market concentration, and greater use of offshoring, particular related to China's entry into the WTO. Assessing the quantitative importance of these alternative mechanisms requires a coherent framework that can be used to fit the actual time series data. Our study represents a small but promising first step in that direction.

Table 1. Calibrated parameters

	Parameter Description	value
β	Subjective discount factor	0.99
ϕ	Unemployment benefit	0.25
α	Elasticity of matching function	0.50
μ	Matching efficiency	0.6629
$\bar{\delta}$	Job separation rate	0.10
ρ^o	Automation obsolescence rate	0.03
κ	Vacancy posting cost	0.0884
b	Nash bargaining weight	0.50
η_v	Elasticity of vacancy creation cost	1
η_a	Elasticity of automation cost	1
κ_a	Flow cost of automated production	0.98
χ	Mean value of preference shock	0.7267
$\bar{Z}$	Mean value of neutral technology shock	1
$\bar{\zeta}$	Mean value of equipment-specific technology shock	1

Table 2. Estimated parameters

		Priors		Posterior		
	Parameter description	Type	[mean, std]	Mean	5%	95%
$\bar{e}$	scale for vacancy creation cost	G	[5, 1]	9.5693	7.4499	11.8171
$\bar{x}$	scale for robot adoption cost	G	[5, 1]	2.4349	1.7760	3.0439
ρ_z	AR(1) of neutral technology shock	B	[0.8, 0.1]	0.9704	0.9616	0.9800
ρ_θ	AR(1) of discount factor shock	B	[0.8, 0.1]	0.9832	0.9717	0.9942
ρ_δ	AR(1) of separation shock	B	[0.8, 0.1]	0.9425	0.9149	0.9677
ρ_ζ	AR(1) of automation-specific shock	B	[0.8, 0.1]	0.7565	0.7201	0.7937
σ_z	std of tech shock	IG	[0.01, 0.1]	0.0111	0.0103	0.0119
σ_θ	std of discount factor shock	IG	[0.01, 0.1]	0.0146	0.0112	0.0180
σ_δ	std of separation shock	IG	[0.01, 0.1]	0.0500	0.0462	0.0537
σ_ζ	std of automation-specific shock	IG	[0.01, 0.1]	0.0423	0.0307	0.0535

Note: This table shows our benchmark estimation results. For the prior distribution types, we use G to denote the gamma distribution, B the beta distribution, and IG the inverse gamma distribution.

Table 3. Forecasting Error Variance Decomposition

Variables	Neutral technology shock	Discount factor shock	Job separation shock	Automation-specific shock
Unemployment	42.84	56.34	0.51	0.32
Vacancy	40.96	53.85	4.76	0.43
Productivity growth	49.02	17.93	0.17	32.88
Real wage growth	26.58	23.99	0.61	48.82
Automation probability	46.93	42.09	0.05	10.94

Note: The numbers reported are the posterior mean contributions (in percentage terms) of each of the four shocks in the benchmark estimation to the forecast error variances of the variables listed in the rows.

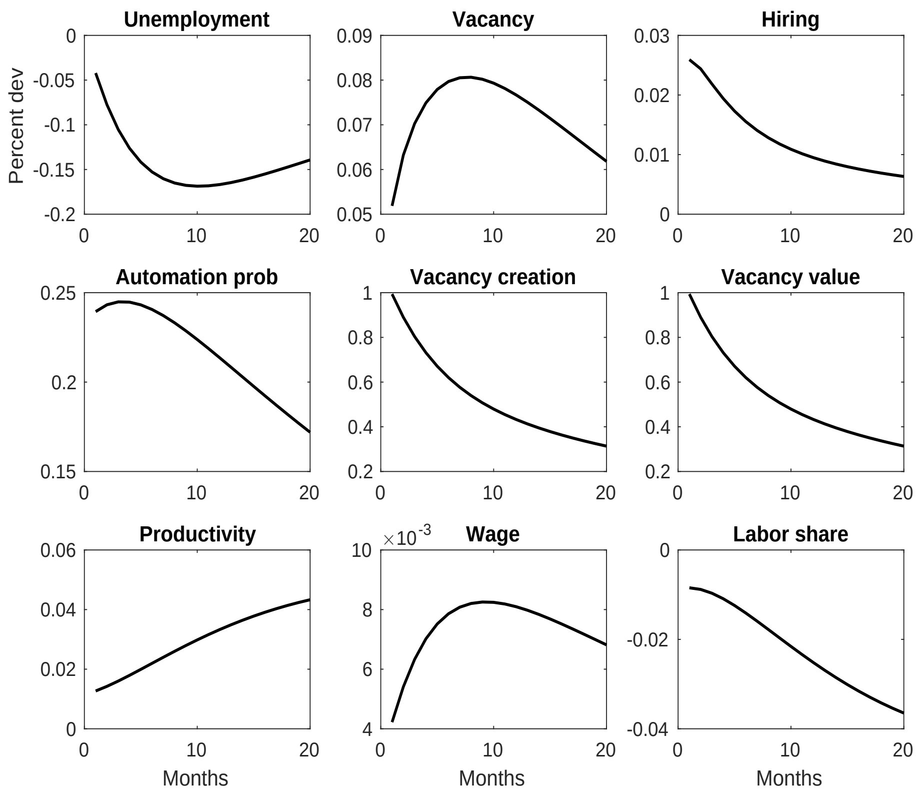

FIGURE 1. Impulse responses to a positive neutral technology shock in the benchmark model.

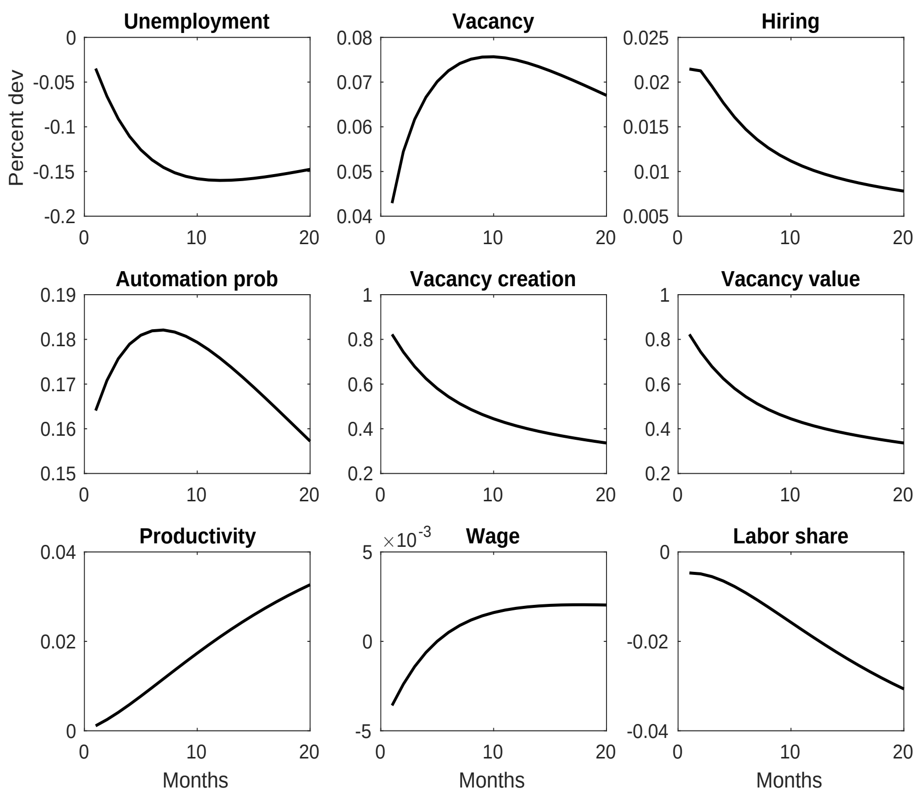

FIGURE 2. Impulse responses to a positive discount factor shock in the benchmark model.

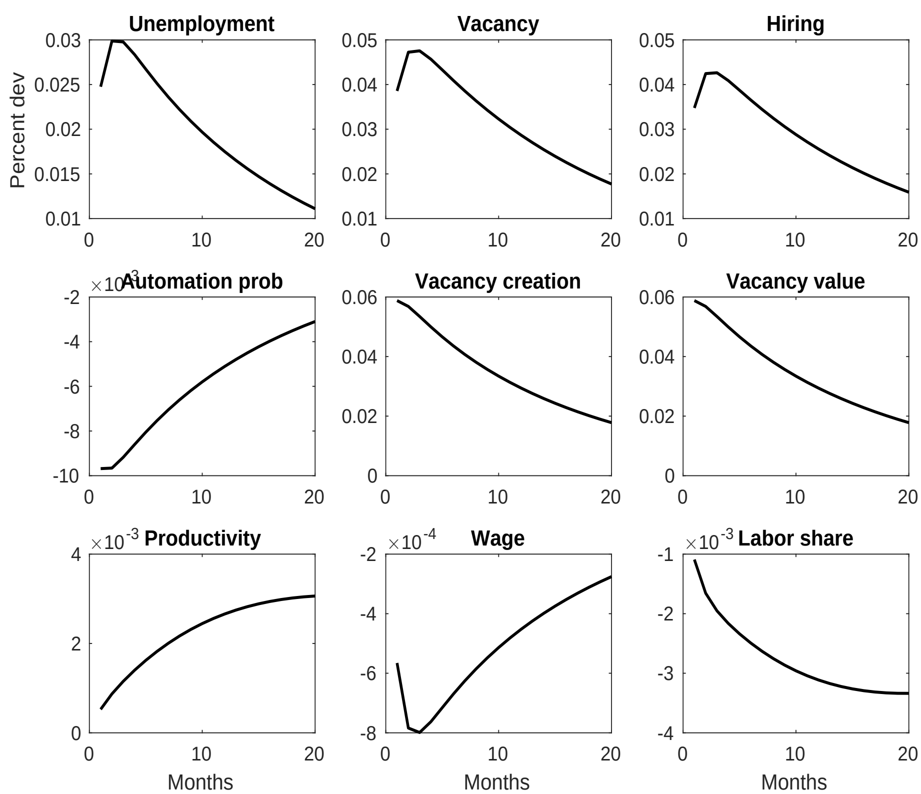

Figure 3. Impulse responses to a job separation shock in the benchmark model.

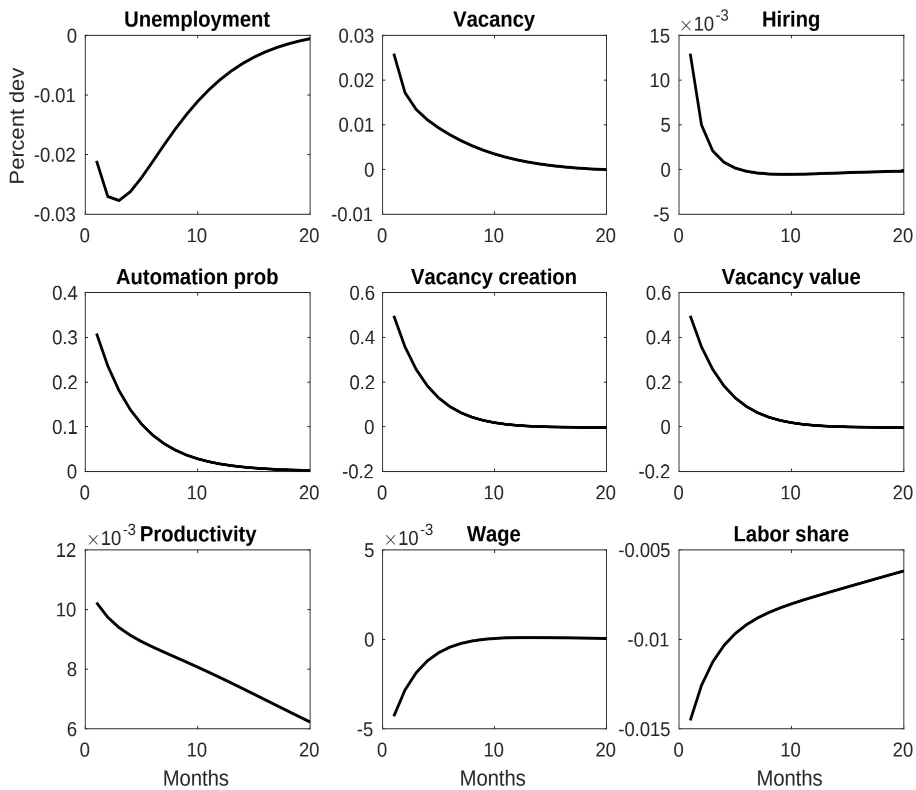

FIGURE 4. Impulse responses to a positive automation-specific shock in the benchmark model.

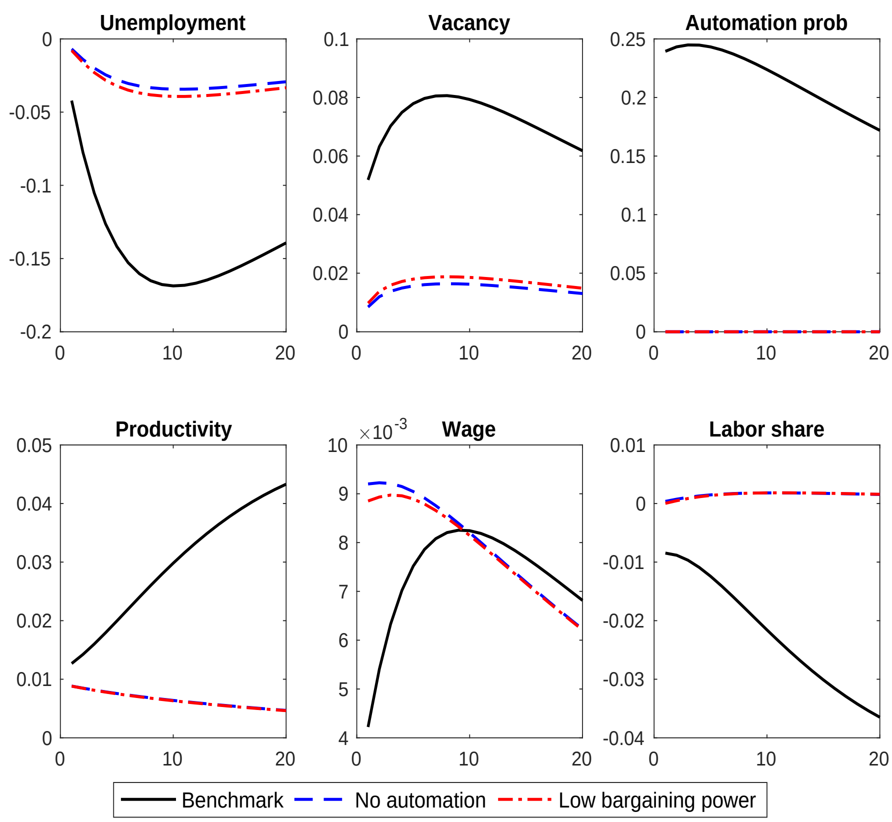

FIGURE 5. Impulse responses to a positive neutral technology shock in the benchmark model (black solid lines), the counterfactual with no automation (blue dashed lines), and the counterfactual with no automation and low worker bargaining power (red dashed and dotted lines).

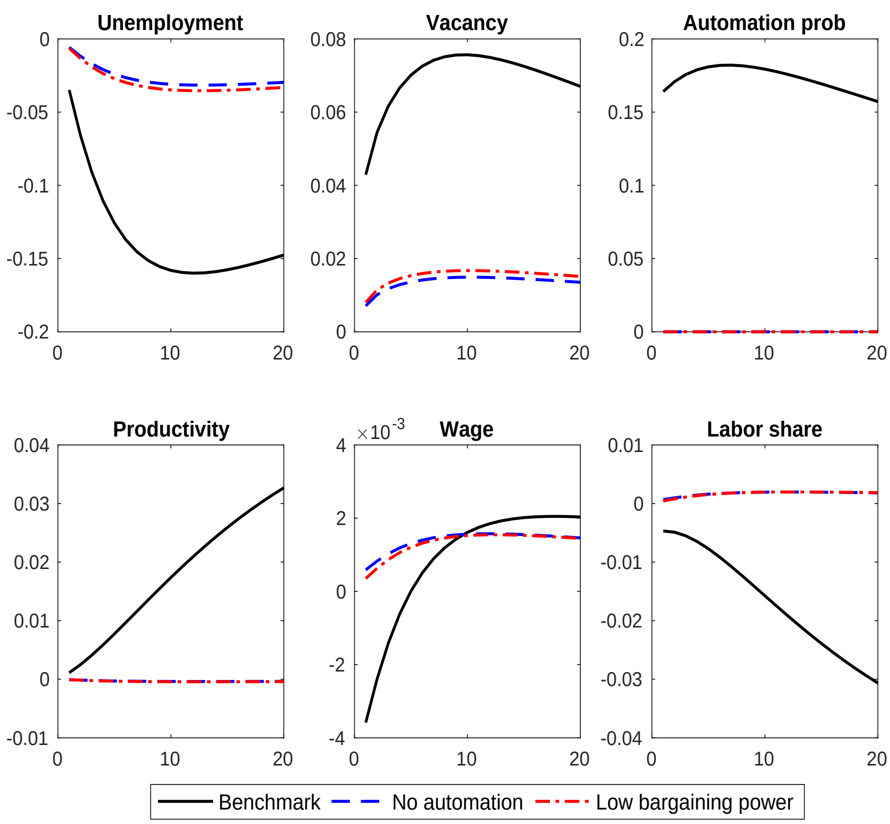

FIGURE 6. Impulse responses to a positive discount factor shock in the benchmark model (black solid lines), the counterfactual with no automation (blue dashed lines), and the counterfactual with no automation and low worker bargaining power (red dashed and dotted lines).

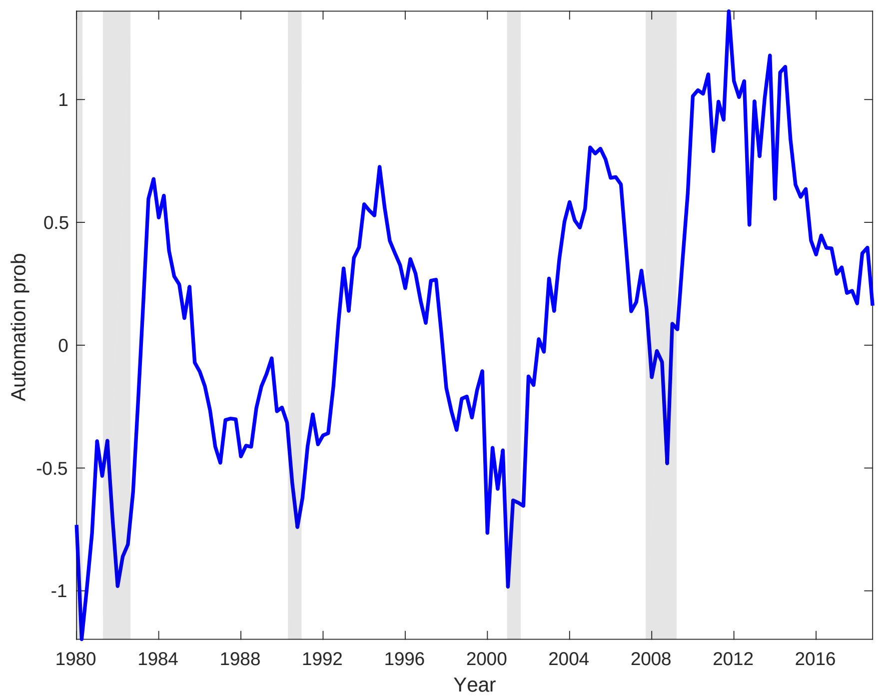

FIGURE 7. The smoothed time series of the automation probability.

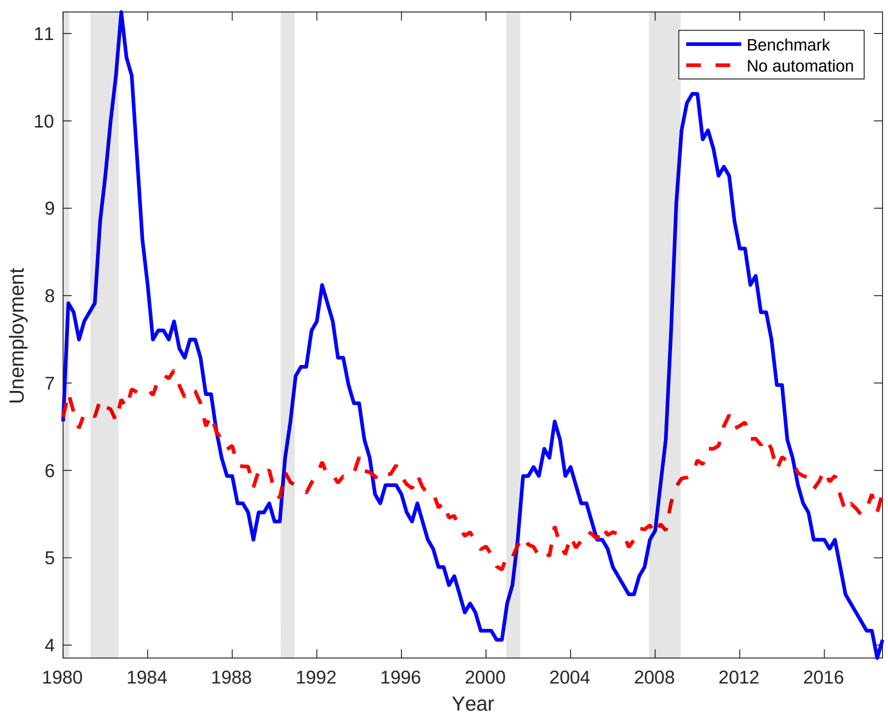

FIGURE 8. The smoothed time series of unemployment in the benchmark model (the blue solid line) versus the counterfactual with no automation (the red dashed line). The counterfactual is a version of our benchmark model with all automation related variables (including automation specific shocks) kept at their steady-state values. The benchmark model series is identical to the actual data because we estimated the model to fit the unemployment series.

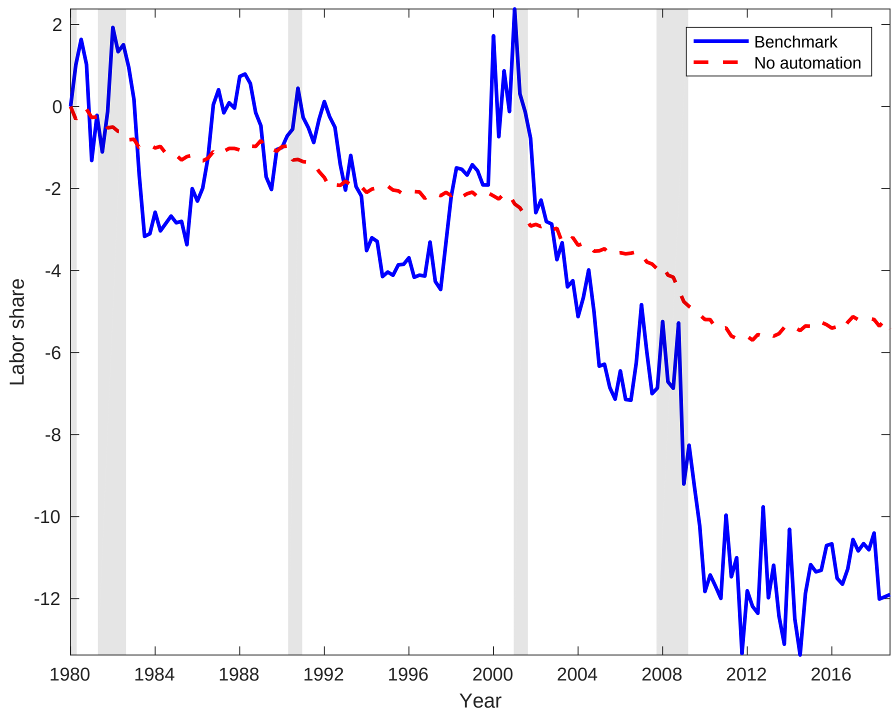

Figure 9. The smoothed time series of the labor share in the benchmark model (the blue solid line) versus the counterfactual with no automation (the red dashed line). The counterfactual is a version of our benchmark model with all automation related variables (including automation specific shocks) kept at their steady-state values. The benchmark model series is identical to the actual data because we estimated the model to fit the real wage and labor productivity series, and the labor share can be directly inferred from these two variables.

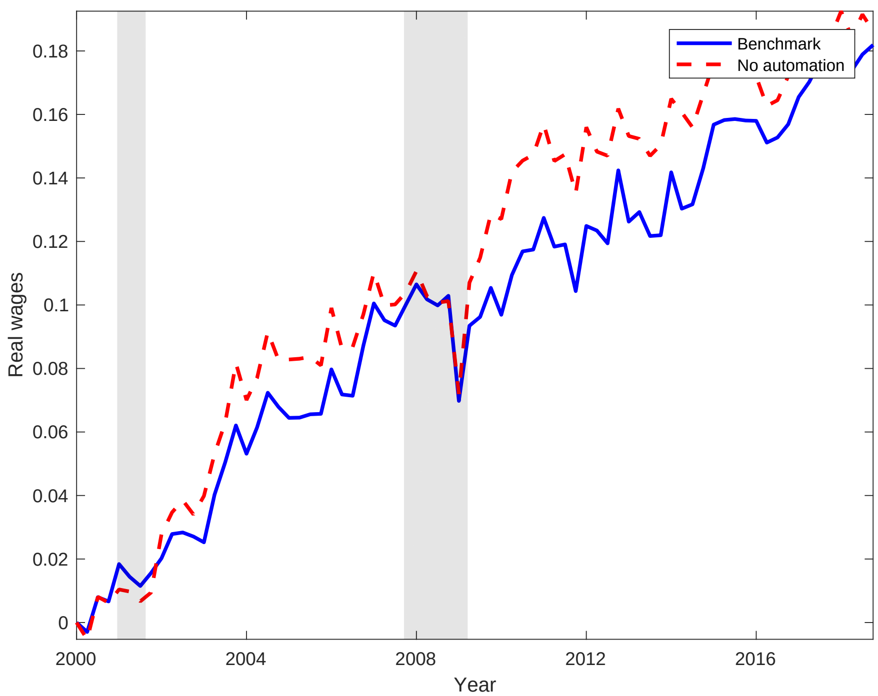

Figure 10. The smoothed time series of real wages in the benchmark model (the blue solid line) versus the counterfactual with no automation (the red dashed line). The counterfactual is a version of our benchmark model with all automation related variables (including automation specific shocks) kept at their steady-state values. The real wage levels are computed based on the smoothed real wage growth series in each model (with the initial value fixed at the actual real wage level in the beginning of the sample).

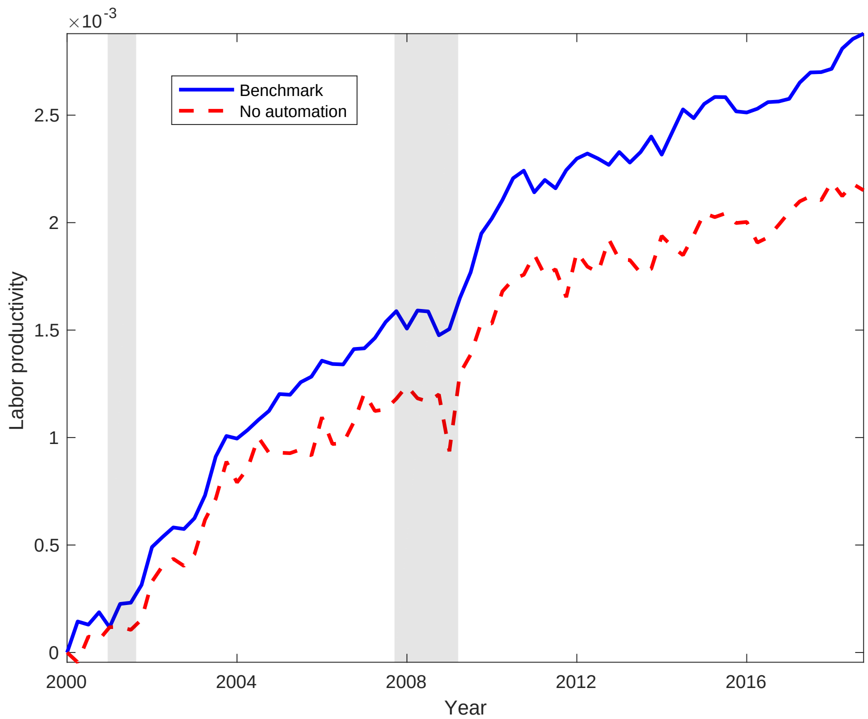

FIGURE 11. The smoothed time series of labor productivity in the benchmark model (the blue solid line) versus the counterfactual with no automation (the red dashed line). The counterfactual is a version of our benchmark model with all automation related variables (including automation specific shocks) kept at their steady-state values. The labor productivity levels are computed based on the smoothed labor productivity growth series in each model (with the initial value fixed at the actual productivity level in the beginning of the sample).

References

Acemoglu, D. and P. Restrepo (2017): "Robots and Jobs: Evidence from US Labor Markets," NBER Working Paper No. 23285.

——— (2018): "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108, 1488–1542.

Arnoud, A. (2018): "Automation Threat and Wage Bargaining," Unpublished Manuscript, Yale University.

Autor, D. H. (2015): "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 29, 3–30.

Autor, D. H. and A. Salomons (2018): "Is automation labor-displacing? Productivity growth, employment, and the labor share," *Brookings Papers on Economic Activity*, 1, 1–87.

Barnichon, R. (2010): "Building a Composite Help-Wanted Index," *Economic Letters*, 109, 175–178.

Blanchard, O. J. and J. Galí (2010): "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment," *American Economic Journal: Macroeconomics*, 2, 1–30.

den Haan, W. J., G. Ramey, and J. Watson (2000): "Job Destruction and Propagation of Shocks," *American Economic Review*, 90, 482–498.

Dinlersoz, E. and Z. Wolf (2018): "Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing," U.S Census Bureau Center for Economic Studies Working Paper 18-39.

Elsby, M. W. L., B. Hobijn, and A. Sahin (2013): "The Decline of the U.S. Labor Share," *Brookings Papers on Economic Activity*, Fall, 1–63.

Fernald, J. G. (2015): "Productivity and Potential Output Before, During, and After the Great Recession," *NBER Macroeconomics Annual 2014*, 29, 1–51.

Frey, C. B. and M. A. Osborne (2015): "The Future of Employment: How Susceptible are Jobs to Computerisation?" *Technological Forecasting and Social Change*, 114, 254—280.

——— (2017): "The Future of Employment: How Susceptible Are Jobs to Compterisation?" *Technological Forecasting and Social Change*, 114, 254–280.

Fujita, S. and G. Ramey (2007): "Job Matching and Propagation," *Journal of Economic Dynamics & Control*, 31, 3671–3698.

Gertler, M. and A. Trigari (2009): "Unemployment Fluctuations with Staggered Nash Wage Bargaining," *Journal of Political Economy*, 117, 38–86.

Graetz, G. and G. Michaels (2018): "Robots at Work," *The Review of Economics and Statistics*, 107, 753–768.

Hall, R. E. (2017): "High Discounts and High Unemployment," *American Economic Review*, 107, 305–330.

Hall, R. E. and P. R. Milgrom (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98, 1653–1674.

Karabarbounis, L. and B. Neiman (2013): "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 129, 61–103.

Krueger, A. B. (2018): "Reflections on Dwindling Worker Bargaining Power and Monetary Policy," Luncheon Address at the Jackson Hole Economic Symposium.

Leduc, S. and Z. Liu (2019): "The Weak Job Recovery in a Macro Model of Search and Recruiting Intensity," *American Economic Journal: Macroeconomics*, (forthcoming).

Nedelkoska, L. and G. Quintini (2018): "Automation, skills use and training," OECD Social, Employment and Migration Working Papers, No. 202.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25—49.

Zeira, J. (1998): "Workers, Machines, and Economic Growth," *Quarterly Journal of Economics*, 113, 1091–1117.

Appendix A. Summary of equilibrium conditions

A search equilibrium is a system of 19 equations for 19 variables summarized in the vector

$$\left[C_t, r_t, Y_t, m_t, u_t, v_t, q_t^u, q_t^v, q_t^a, N_t, U_t, \eta_t, J_t^e, J_t^v, J_t^a, A_t, x_t^*, w_t^N, w_t\right].$$

We write the equations in the same order as in the dynare code.

(1) Household's bond Euler equation:

$$1 = \mathrm{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} r_t, \tag{A1}$$

(2) Matching function

$$m_t = \mu u_t^{\alpha} v_t^{1-\alpha}, \tag{A2}$$

(3) Job finding rate

$$q_t^u = \frac{m_t}{u_t}, \tag{A3}$$

(4) Vacancy filling rate

$$q_t^v = \frac{m_t}{v_t}, \tag{A4}$$

(5) Employment dynamics:

$$N_t = (1-\delta_t) N_{t-1} + m_t, \tag{A5}$$

(6) Number of searching workers:

$$u_t = 1 - (1-\delta_t) N_{t-1}, \tag{A6}$$

(7) Unemployment:

$$U_t = 1 - N_t, \tag{A7}$$

(8) Vacancy dynamics

$$v_t = (1 - q_{t-1}^v)(1 - q_{t-1}^a) v_{t-1} + \delta_t N_{t-1} + \eta_t, \tag{A8}$$

(9) Automation dynamics

$$A_t = (1-\rho^o) A_{t-1} + (1 - q_t^v) q_t^a v_t, \tag{A9}$$

(10) Employment value

$$J_t^e = Z_t - w_t + \mathbb{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} \left\{ (1-\delta_{t+1}) J_{t+1}^e + \delta_{t+1} J_{t+1}^v \right\}, \tag{A10}$$

(11) Vacancy value

$$J_t^v = -\kappa + q_t^v J_t^e + (1-q_t^v) q_t^a J_t^a + (1-q_t^v)(1-q_t^a) \mathbb{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} J_{t+1}^v. \tag{A11}$$

(12) Automation value

$$J_t^a = Z_t\zeta_t - \kappa_{at} + (1-\rho^o)\mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}J_{t+1}^a, \tag{A12}$$

(13) Automation threshold

$$x_t^* = J_t^a - \mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}J_{t+1}^v. \tag{A13}$$

(14) Robot adoption

$$q_t^a = \left(\frac{x_t^*}{\bar{x}}\right)^{\eta_a} \tag{A14}$$

(15) Vacancy creation

$$\eta = \frac{J_t^v}{\bar{e}} \tag{A15}$$

(16) Aggregate output

$$Y_t = Z_tN_t + Z_t\zeta_tA_t. \tag{A16}$$

(17) Resource constraint

$$C_t + \kappa v_t + \frac{\eta_a}{1+\eta_a}q_t^a x_t^*(1-q_t^v)v_t + \frac{1}{2}\eta_t J_t^v = Y_t, \tag{A17}$$

(18) Nash bargaining wage:

$$\frac{b}{1-b}(J_t^e - J_t^v) = w_t^N - \phi - \chi C_t + \mathbb{E}_t\frac{\beta\theta_{t+1}C_t}{C_{t+1}}(1-q_{t+1}^u)(1-\delta_{t+1})\frac{b}{1-b}(J_{t+1}^e - J_{t+1}^v). \tag{A18}$$

(19) Actual real wage (with real wage rigidity)

$$w_t = w_{t-1}^{\gamma}(w_t^N)^{1-\gamma}, \tag{A19}$$